

Temporal Opinion Spam Detection by Multivariate Indicative Signals

Junting Ye

Santhosh Kumar

Leman Akoglu

Stony Brook University
Department of Computer Science
{juyye,smanavasilak,leman}@cs.stonybrook.edu

Abstract

Online consumer reviews reflect the testimonials of real people, unlike advertisements. As such, they have critical impact on potential consumers, and indirectly on businesses. According to a Harvard study (Luca 2011), +1 rise in star-rating increases revenue by 5–9%. Problematically, such financial incentives have created a market for spammers to fabricate reviews, to unjustly promote or demote businesses, activities known as opinion spam (Jindal and Liu 2008). A vast majority of existing work on this problem have formulations based on *static* review data, with respective techniques operating in an *offline* fashion. Spam campaigns, however, are intended to make most impact during their course. Abnormal events triggered by spammers’ activities could be masked in the load of future events, which static analysis would fail to identify. In this work, we approach the opinion spam problem with a *temporal* formulation. Specifically, we monitor a list of carefully selected indicative signals of opinion spam over time and design *efficient* techniques to both detect and characterize abnormal events in *real-time*. Experiments on datasets from two different review sites show that our approach is fast, effective, and practical to be deployed in real-world systems.

Introduction

Online product reviews play important role for e-commerce. New customers tend to prefer products with higher ratings as previous buyers have “testified” that the products are good choices. On the contrary, new customers may have less interest for products with lower ratings since more dissatisfaction from experienced users have been reported. Driven by such commercial benefits, spam or fake reviews have become a prevalent problem, for which effective detection algorithms are greatly needed (Jindal and Liu 2008).

In the past years, several existing works employed supervised techniques by extracting features based on review text, ratings, product meta-data such as category and price, review feedback such as number of helpful votes, etc. (Jindal and Liu 2008; Ott et al. 2011; Feng, Banerjee, and Choi 2012; Mukherjee et al. 2013a). These methods have a key challenge as ground truth is extremely hard to obtain—human annotators can hardly tell the difference between the genuine and fake reviews (Ott et al. 2011).

Unsupervised approaches have also been explored for the opinion spam problem. In a nutshell, these can be categorized as approaches that leverage linguistic (Ott et al. 2011; Feng, Banerjee, and Choi 2012), relational (Wang et al. 2011; Akoglu, Chandy, and Faloutsos 2013; Li et al. 2014; Ye and Akoglu 2015), and most often, behavioral clues (Jindal, Liu, and Lim 2010; Feng et al. 2012; Mukherjee, Liu, and Glance 2012; Lim et al. 2010; Mukherjee et al. 2013a; Xie et al. 2012). In their investigation, (Mukherjee et al. 2013b) found that Yelp’s fake review filter might be relying more on behavioral rather than text-based clues. Most recently, (Rayana and Akoglu 2015) proposed a unifying approach that harnesses all of these information sources, which outperforms any individual one alone, especially when provided with a few labeled instances.

As one can see, previous work has focused on aspects such as availability of labels or the lack thereof, or the type of information to consume (linguistic, relational, behavioral). Orthogonal to these aspects, in this work we bring emphasis on the aspect of *time*. Surprisingly, we realize that the vast majority of past work has formulated the opinion spam problem on *static* review data, and developed techniques that operate in an *offline* fashion. A few works, such as those by (Fei et al. 2013) and (Li et al. 2015), designed and used temporal features along with static ones. Nevertheless, their approaches are effectively offline where the entire review data is assumed to be available at analysis time. Recently, (Xie et al. 2012) proposed an approach that tracks three temporal features—average rating, number of reviews, and ratio of singleton reviewers—over time, and pinpoints time windows in which all three change simultaneously. Their approach is specifically crafted for catching *singletons*, i.e., spam reviewers with only a single review, that earlier methods mainly failed to effectively identify.

In this work, we consider opinion spam as a *temporal* phenomenon. Our intuition is that the spammers’ activities trigger abrupt changes and unusual events in the underlying generating process *at the time* of their operations. Such changes may later be reverted, or the impact of such events may be lost in the abundance of future events. As such, we argue that a temporal, *real-time* approach would not only be more effective but also more realistic, as review data is temporal in nature. In the following we provide an overview of our proposed methodology, and the list of contributions.

Overview

To promote or demote products, spammers need to dominate the sentiments among a given product’s reviews, which requires considerable amount of reviews. Such intensive behaviors result in abnormal bursts of various suspicious activities over the timeline of the product. In this work, our goal is twofold: we aim to *both* detect *and* characterize those abnormal bursts of activities.

In a nutshell, we first identify a number of signals associated with suspicious reviewing behavior. For each product (or business), these signals are tracked over time in an online fashion. Through efficient time series modeling, we detect anomalous time points among all the signals. Finally, we leverage the specific anomalous signals for characterization, i.e., inspection and sensemaking. We outline the main steps of our approach as follows:

1. *Temporal signal extraction*: We define a list of measures, called *indicative signals*, that provide potential evidence toward opinion spam. These measures are computed for each product over time as new reviews arrive. As such, each product is represented by multiple time series.
2. *Anomaly detection in lead (target) signal*: We designate one of the signals as *the lead*. Anomalous changes in the lead signal suggest potential spamming activity.
3. *Anomaly detection in supporting signals*: Anomalies in the lead signal provide the necessary conditions for opinion spam; but not necessarily sufficient. For endorsement, we perform anomaly detection on the remaining signals in an efficient way, by focusing on time periods around the anomalous points (“alarms”) detected in the lead.
4. *Characterization*: Our method spots anomalous time points across all signals for each product. We use a function of the number and magnitude of the anomalies to rank the products. At any given time, products are investigated in their ranked order. The specific detected time points and support signals guide the manual inspection.

Contributions

- **Problem formulation**: We propose a new temporal approach to the opinion spam problem, that monitors a carefully selected list of indicative signals. Anomalies across various signals suggest and corroborate evidence. As shown in the experiments, our approach is (i) *descriptive*; it facilitates inspection and sensemaking, and (ii) *general*; it can identify different kinds of spam campaigns.
- **A new methodology**: We develop an *online* and *efficient* algorithm for opinion spam. We mainly monitor the lead signal and detect anomalies in *real-time*. Only when an “alarm” is raised in the lead, a *local* algorithm then performs anomaly detection in the support signals by considering a small window around the anomaly (hence local).
- **Practicality**: The proposed approach is usable in practice, as it provides capabilities for online monitoring, automatic detection, and interpretability for the analysts. Specifically, detected time points narrow down where to look, and the indicative signals suggest what to look at. We demonstrate the practicality of our approach through case studies in two review datasets, for apps in an online marketplace and for products in Flipkart (cf. Figure 1).




Figure 1: Time series for 9 indicative signals for a software product (app) (green curves: signal values, red curves: anomaly score, blue bars: anomalous time points). Our method detects 4 weeks of attacks that increase the average rating of a product with generally declining rating. Results discussed further in experiments.

Reproducibility: Open source code of our method is available at: http://www3.cs.stonybrook.edu/~juyye/code/ICWSM16_Code.zip.

Indicative Signals of Opinion Spam

We consider the review data of a product p as a stream of review units ordered in time, $D_p = \{u_1, u_2, \dots, u_i, \dots\}$. Each unit consists of user ID, timestamp, review text or content, and rating; i.e, $u_i = (UID_i, t_i, c_i, r_i)$, where $t_i \leq t_j$ if $i < j$. Timestamps are divided into fixed-length time windows (ΔT) where the values of each time series are computed based on data from these intervals.

We identify eight types of indicative signals, and extract their time series for each product p . A list of these signals are given in Table 1 which we describe as follows.

1. **Average Rating**: If the average rating of a product p changes significantly, it might have been tempered with by spammers. This time series tracks the evolution of *cumulative* average rating. Let $\bar{U}_p^t = \{u_k | t_k \in [0, t * \Delta T)\}$ denote the set of p ’s reviews until the end of time window t , where $|\bar{U}_p^t| = m$. Then,

$$R_p^t = \frac{1}{m} \sum_{k=1}^m r_k \quad (1)$$

Table 1: Indicative signals of opinion spam.

Name	Range	Suspicious if
Average Rating	[1, 5]	Change
Number of (+/-) Reviews	[0, ∞]	Increase
Rating Entropy	[0, $\log_2 5$]	Decrease
Ratio of Singletons	[0, 1]	Increase
Ratio of First-timers	[0, 1]	Increase
Youth Score	[0, 1]	Increase
Temporal Gap Entropy	[0, \max_e^\dagger]	Decrease

\dagger With windows size ΔT and logarithmic binning, number of bins is $\lceil \log_2 \Delta T \rceil + 1$ and $\max_e = \log_2(\lceil \log_2 \Delta T \rceil + 1)$.

2. **Number of Reviews:** A sudden surge in the number of reviews may indicate spamming activity. As such, we track the total number of reviews within each time interval. Let $U_p^t = \{u_k | t_k \in [(t-1) * \Delta T, t * \Delta T]\}$ denote p 's reviews within window t . Then,

$$C_p^t = |U_p^t| \quad (2)$$

3. **Number of Positive/Negative Reviews:** In addition to total number of reviews, we also track the positive and negative review counts, as fake reviews either demote or promote and have skewed ratings.

$$+C_p^t = |\{u_k | u_k \in U_p^t, r_k \in \{4, 5\}\}| \quad (3)$$

$$-C_p^t = |\{u_k | u_k \in U_p^t, r_k \in \{1, 2\}\}| \quad (4)$$

4. **Rating Entropy:** We also monitor entropy over time, as a measure of skewness in ratings. As such,

$$E_p^t = - \sum_{r=1}^5 p_r^t \cdot \log p_r^t \quad (5)$$

where p_r^t is the ratio of reviews with rating value equal to r in window t .

5. **Ratio of Singletons:** Fake reviews could be posted by newly created accounts. Therefore, we track the ratio of one-time reviewers over time.

$$S_p^t = \frac{|U_s^t|}{C_p^t} \quad (6)$$

where U_s^t is the set of singleton users who posted their first-ever *and only* review (to p) during window t .

6. **Ratio of First-timers:** Some campaigns involve spammers targeting multiple products simultaneously. As such, we also track the ratio of first-time reviewers as

$$F_p^t = \frac{|U_f^t|}{C_p^t} \quad (7)$$

where U_f^t is the set of p 's reviewers who posted their first but not necessarily only review during window t (note that in this case the first review need not be to p).

7. **Youth Score:** Most fake reviews are posted from short-lived accounts. Therefore, besides singletons and first-timers, we further track account age. Specifically, for each review u_k of p within window t , we compute the age of its reviewer UID_k at the time they posted u_k by

$$A_k = t_k - t_0^{UID_k}$$

where $t_0^{UID_k}$ is the time at which reviewer UID_k posted their first review. The youth score is then the average of reviewer ages at the time they posted for p .

$$Y_p^t = \frac{1}{C_p^t} \sum_{u_k \in U_p^t} 2 \cdot \left(1 - \frac{1}{1 + \exp(-A_k)}\right) \quad (8)$$

The larger the A_k 's, the closer Y_p^t gets to zero. As such, larger Y_p^t is more suspicious.

8. **Temporal Gap Entropy:** A normal user is expected to post reviews at arbitrary intervals, while spammers may follow a predictable robot-like behavior (e.g., post every day). As such, we compute the time-gap between consecutive reviews $\{u_k, u_{k+1}\} \in U_p^t$ of p within window t , create a histogram, and compute the entropy. The histogram is constructed through logarithmic binning (Newman 2005) (e.g., if ΔT is 30 days, consecutive bin sizes are 0, 1, 2, 4, 8, and 16 days). That is,

$$G_p^t = - \sum_{b=1}^{\lceil \log_2 \Delta T \rceil + 1} p_b^t \cdot \log p_b^t \quad (9)$$

where p_b^t is the ratio of time-gaps in bin b in window t .

Notice that the temporal signals of interest are potential indicators of spamming behavior. However, they need not exhibit abnormalities altogether at the same time windows. In fact, different spam campaigns might trigger different signals to fire. For example, one campaign may create a large number of new accounts but use them in arbitrary time intervals. Another may involve a very skewed number of positive ratings but from sufficiently old accounts. Our approach is multi-faceted, as such it could help identify various combinations of collusions among spammers.

Temporal Opinion Spam Detection

Our temporal approach to opinion spam detection consists of four main steps: (i) extracting temporal signal values as defined in the previous section; (ii) detecting changes in what is called the *lead* signal; (iii) checking whether we also observe temporal anomalies in the supporting signals; and (iv) ranking targeted products based on the number and magnitude of anomalies found in their timeline.

In the following subsections, we describe the details of the last three steps, respectively.

Anomalies in the Lead (Target) Signal

Out of all the indicative signals that we extract over time, we dedicate one of them as the *lead* signal. The lead can be chosen as the measure that spammers particularly aim to manipulate, e.g., the average rating. It can also be a measure for which spamming activities could trigger a change, e.g., the number of (positive, negative, or total) reviews.

To elaborate, average rating serves as the overall impression of consumers for a given product. It has the potential to shape the first impression of candidate consumers. Therefore, many spammers are devoted to manipulate the average rating of products (Lim et al. 2010). If there is a sudden change in average rating over time (either increase or decrease), it is possible that the product is a target of a (promotion or demotion) spam campaign.

On the other hand, it gets harder and harder to change the average rating for products with increasingly many reviews, compared to those with only a few. In their analysis, (Rahman et al. 2015) show that the minimum number of reviews needed to fraudulently increase the average rating of a product by 1/2 star is $n/7$, where n is the number of genuine reviews of the product. This suggests that it is costlier to change the average rating of popular (i.e., famous or notorious) products with large number of reviews, provided spammers are paid per review (Streitfeld 2011). Nevertheless, as we show in the experiments, there exist scenarios where we observe a burst of fake reviews in the timeline of a product, even though those do not have any impact on the average rating. The incentives can be various; such as promoting a new feature of a product, flooding the most recent page view (as reviews are often ordered in time), etc. In order to spot such spam campaigns, we use number of positive (or negative) reviews as the lead signal and watch for abnormal increases.

For anomaly detection on the lead time series, one can use any detection algorithm that provides real-time capability. One important issue is the semantics of the lead signal. Average rating is cumulative, and our goal is to find *change points* in both directions (either increase or decrease). For this lead, we use the cumulative sum (CUSUM) change detection algorithm (Page 1954). On the other hand, we track the non-cumulative number of (+ or -) reviews per time window ΔT , with a goal to spot *anomalies* in the form of bursts (i.e., large increases). For such leads, we use the autoregressive (AR) model for modeling the lead series, and use the deviation from the forecasted value as the anomaly score s at a new time point.

For change/anomaly detection, choice of a threshold is critical to flag alerts. For a given (lead) signal and a detector, we maintain the distribution $\mathcal{D}(S|T, P)$ of the anomalousness scores $S(i)$ across time points T and (ii) across products P . We then employ Cantelli’s inequality¹ to identify a theoretical threshold $\delta = (\mathcal{D}, \eta)$, where η is the expected percentage of anomalies (see (Smets and Vreeken 2011)). For a given score s_p^t for product p at time t , we flag an alert if $s_p^t > \delta$ and the anomaly is in the direction of what a suspicious activity would create (e.g., increase in youth score but decrease in entropy, see Table 1).

Anomalies in the Supporting Signals

Lead signal alone is not sufficient to indicate the occurrence of spamming activities. For example, when a sudden increase is detected in lead signal, e.g. average rating, the causes could be (i) the product owner hired spammers to write fake reviews for promotion; (ii) the product’s quality has improved such that genuine reviewers wrote more positive reviews. Therefore, we investigate further the supporting signals, to verify if “alarms” triggered by the lead signal are indeed the consequences of spamming activities.

In order to detect anomalies in each supporting signal, we first propose GLOBALAR, which is an *online* algorithm that quantifies the extent an input value v_i is an anomaly in time

Table 2: *Notation used throughout text.*

Notation	Definition
v_{i-k}^{i-1}	a sequence of input values $(v_{i-k}, \dots, v_{i-1})$
t_i	i_{th} time point
t_a	most recent alarm time point
k	order of AR (autoregressive) model
r	discounting parameter of SDAR
δ	anomalousness score threshold
ω	AR coefficients
μ	mean of time series values
σ	variance of time series values
θ	$\theta = \{\omega, \mu, \sigma\}$
s_{i-4}^{i-1}	a sequence of square errors $(s_{i-4}, \dots, s_{i-1})$
O_{i-4}^{i-1}	a sequence of anomaly labels $(O_{i-4}, \dots, O_{i-1})$
L	number of data points used to pick k

series setting. Its major component is the Sequentially Discounting Auto Regression algorithm (SDAR) (Yamanishi and Takeuchi 2002). SDAR detects temporal anomalies by modeling a given time series with autoregressive model and discounting the importance of historical data exponentially.

GLOBALAR is effective but computationally expensive as it computes the anomaly score for each and every time point in the supporting signals. Therefore, we modify GLOBALAR and propose LOCALAR to reduce time complexity. Given a supporting signal, LOCALAR only focuses on and scores the time points around the “alarms” produced by the lead signal, and hence selectively ignores the other time points. This reduces the time complexity from linear to *sublinear*, in terms of the total number of time points. In the following subsections, we elaborate on the technical details of GLOBALAR and LOCALAR, respectively.

GLOBALAR GLOBALAR detects temporal anomalies by using SDAR. Simply put, SDAR predicts value (denoted as \hat{v}_i) at time point t_i and computes square error between prediction and observation, i.e., $(v_i - \hat{v}_i)^2$. If the square error is large, then v_i is more likely to be an anomaly. It is easy to see that the square error can be large when \hat{v}_i is very large or very small. However, in the setting of opinion spam detection, only either one of the cases indicates the occurrence of spamming activities. For example, a product might have been attacked by spammers if the rating entropy drops significantly (i.e. reviewers during the same time point give very similar ratings). In contrast, if rating entropy suddenly increases, it does not necessarily indicate the activities of spammers. In the remainder of text, we call v_i as “semantically suspicious” (denoted as $SemSus(v_i) = True$) if (1) $v_i > v_{i-1}$ and large value indicates spamming activities (e.g., ratio of singletons), or (2) if $v_i < v_{i-1}$ and small value indicates suspicious behaviors (e.g., rating entropy). Otherwise $SemSus(v_i) = False$.

Next, we introduce the general idea of SDAR. It first models time series $V = \{v_k, v_{k+1}, \dots, v_n\}$ with k_{th} order autoregressive model as in Equation 10 and 11.

$$z_i = \omega z_{i-k}^{i-1} + \epsilon \quad (10)$$

$$v_i = z_i + \mu \quad (11)$$

where $z_{i-k}^{i-1} = (z_{i-k}, \dots, z_{i-1})$, $\omega = (\omega_1, \dots, \omega_k)$, ϵ is a normal random variable with zero mean and σ variance, i.e.

¹Unlike the t-test which assumes Gaussian data, Cantelli’s inequality does not make any distributional assumptions on \mathcal{D} .

Algorithm 1: GLOBALAR

```

1 Input:  $t_i, v_{i-k}, \theta_{i-1}, k, r, t_a, \delta$ 
2 Output:  $\theta_i, s_{t_a-2}^i, O_{t_a-2}^i$ 
3  $\theta_i = \text{SDAR}(v_i, \theta_{i-1}, r, k)$  (where  $\theta_i = \{\omega_i, \mu_i, \sigma_i\}$ )
4  $\hat{v}_i = \omega_i(v_{i-k}^{i-1} - \mu_i) + \mu_i$ 
5  $s_i = (v_i - \hat{v}_i)^2$ 
6 if  $t_i - t_a \leq 2$  then // skip if no recent anomaly in lead signal
7   foreach  $t_j \in [t_a - 2, t_i]$  do // check time points around  $t_a$ 
8     if  $s_j > \delta$  &  $\text{SemSus}(v_j)$  is True then
9        $O_j = 1$ 

```

$\epsilon \sim \mathcal{N}(0, \sigma)$, μ is the mean value of the time series.

Then the probability density function of v_i is defined as Equation 12.

$$p(v_i | v_{i-k}^{i-1}, \theta) = \frac{1}{(2\pi)^{1/2}\sigma} \exp\left(-\frac{(v_i - \hat{v}_i)^2}{2\sigma^2}\right) \quad (12)$$

where $\hat{v}_i = \omega(v_{i-k}^{i-1} - \mu) + \mu$, $v_{i-k}^{i-1} = (v_{i-k}, \dots, v_{i-1})$, μ is a k -dimensional vector with all elements equal to μ , and $\theta = \{\omega, \mu, \sigma\}$.

SDAR discounts the importance of historical data and estimates θ by maximizing the exponentially weighted log likelihood as shown in Equation 13.

$$L(\theta | v_1^i, r, k) = \sum_{m=k+1}^i (1-r)^{i-m} \log p(v_m | v_{m-k}^{m-1}, \theta) \quad (13)$$

where $r \in [0, 1)$ is the discounting factor.

In Equation 13, the estimation of θ is defined on the entire time series. Fortunately, it can also be estimated incrementally. Due to limited space, we skip the technical details of the incremental estimation (See (Yamanishi and Takeuchi 2002) for details). We use $\theta_i = \text{SDAR}(v_i, \theta_{i-1}, r, k)$ to denote the update of parameters from θ_{i-1} to θ_i .

We show the steps of GLOBALAR in Algorithm 1. We first estimate the parameters of SDAR, given a new input value v_i (Line 3). Note that only r and k are hyperparameters. θ_{i-1} is the parameters estimated in the last iteration. Then we predict \hat{v}_i using the estimated parameters θ_i (Line 4). If v_i is significantly different from \hat{v}_i (measured by square error s_i), then v_i is deemed suspicious compared to past data (Line 5). In Lines 6 to 9, we first check if the current time point t_i is close to the most recent “alarm” t_a from the lead signal. If so, we investigate the time points around t_a , particularly $[t_a - 2, t_i]$, instead of simply looking at t_i . The reason is that there might be lags between the lead and the supporting signals. For every time point $t_j \in [t_a - 2, t_i]$, the anomaly label O_j is assigned 1 if the corresponding square error s_j is larger than threshold δ and $\text{SemSus}(v_j)$ returns *True*.

Local AR In GLOBALAR, we investigate all the values in a time series. However, we expect and find that anomalies in the lead signal are much fewer compared to the number of time points. Recall that GLOBALAR exponentially discounts the importance of historical data. It implies that at time point t_i , the values that are much earlier than t_i make little impact when estimating the parameters θ_i . Since we

Algorithm 2: LOCALAR

```

1 Input:  $t_i, t_a, L, v_{i-L-5}^i, \delta$ 
2 Output:  $s_{t_a-2}^i, O_{t_a-2}^i$ 
3 if  $t_i - t_a > 2$  then // exit if no recent anomaly in lead signal
4   Exit
5 foreach  $k' \in [1, 5]$  do // select  $k$  that minimizes square error
6   Init  $S_{k'} = 0$ 
7   foreach  $j \in [1, L]$  do
8      $\theta = \text{AR}(v_{i-j-k'}^{i-j-1}, k')$  (where  $\theta = \{\omega, \mu, \sigma\}$ )
9      $\hat{v}_{i-j} = \omega(v_{i-j-k'}^{i-j-1} - \mu) + \mu$ 
10     $s_{i-j} = (v_{i-j} - \hat{v}_{i-j})^2$ 
11     $S_{k'} = S_{k'} + s_{i-j}$ 
12  $k = k'_{\min}$ , where  $\forall k' \in [1, 5], S_{k'} \geq S_{k'_{\min}}$ 
13 foreach  $t_j \in [t_a - 2, t_i]$  do // check time points around  $t_a$ 
14    $\theta_j = \text{AR}(v_{j-k}^{j-1}, k)$ 
15    $\hat{v}_j = \omega_j(v_{j-k}^{j-1} - \mu_j) + \mu_j$ 
16    $s_j = (v_j - \hat{v}_j)^2$ 
17   if  $s_j > \delta$  &  $\text{SemSus}(v_j)$  is True then
18      $O_j = 1$ 

```

are mainly interested in time points “alarmed” by the lead signal, the question becomes if it would influence the performance much if we only focused on time points close to the “alarms”. Motivated by this, we propose LOCALAR. As we show in our experiments, both algorithms perform similarly but LOCALAR is significantly faster than GLOBALAR.

LOCALAR detects anomalies by using the autoregressive (AR) model. Different from SDAR, AR estimates parameters θ by maximizing the data log likelihood $p(v_i | v_{i-k}^{i-1}, \theta)$ (see Equation 12). The data likelihood is dependent only on v_{i-k}^{i-1} , such that we can ignore values that are far away from v_i . Besides the improvement of efficiency, we pick a proper order, i.e. k , for AR by fitting regular (non-anomalous) values in the time series. We expect that estimating k would be more effective than simply fixing k , as generally, the behaviors of time series from different products are not the same. Similar to GLOBALAR, we allow lags between the lead and supporting signals in LOCALAR.

Algorithm 2 shows the detailed steps of LOCALAR. At time t_i , we check if there is a close-by “alarm” t_a in the lead signal. If not, then we exit the algorithm (Lines 3-4). This step accelerates the algorithm significantly in two aspects: (i) it skips anomaly score computation for points away from the lead “alarms”, as a result of which (ii) feature extraction for a large body of time points in supporting signals can also be skipped. Next, we select a proper (integer) k that minimizes the total square error over a window of L values² before v_i (Lines 5-12). Specifically, we pick a candidate k' and initialize the square error sum $S_{k'}$ to 0. We compute the square error s_{i-j} between the inputs v_{i-j} and predictions \hat{v}_{i-j} of an AR model of order k' , for all L values before v_i . Sum of square errors is denoted by $S_{k'}$. We then choose the k' with the minimum $S_{k'}$ as the order of our AR model at

²Window size L is essentially the training data size used to estimate k . Choice of L poses a trade-off between the estimation quality and running time. In experiments we find $L = 8$ effective.

time t_i . As temporal dependence drops by distance in time, we focus on small $k' \in [1, 5]$. Through Lines 13-18, we carefully examine the time points around t_a . Different from GLOBALAR, we compute the square error for values at and before v_i at this step, using the estimated k . When the square error is larger than the anomaly threshold and $SemSus(v_j)$ returns *True*, anomaly label O_j is assigned 1.

Scoring and Ranking Products

For a product p_i at time point t_j , we can use CUSUM, GLOBALAR or LOCALAR to detect anomalies in both its lead and supporting signals. Given these anomalies, how can we quantify the suspiciousness of p_i at t_j ? How can we rank p_i among all other products at that time?

We answer these questions by formulating a suspiciousness score for the products. Intuitively, it is more probable that a product is a target of spamming activities (i) if there are a large number of temporal anomalies among its indicative signals, and (ii) if the magnitudes of the anomalies are large. Based on these insights, we design four measures to quantify product suspiciousness.

First is the ratio of anomalies among product p_i 's nine indicative signals at t_j . That is, $f_1(p_i, t_j) = \sum_{l=1}^9 O_{j,p_i}^{(l)} / 9$, where $O_{j,p_i}^{(l)} \in \{0, 1\}$ is the anomaly label of signal l of product p_i at t_j . The second and third measures are respectively the average and maximum magnitude of the anomalies, and can be written as $f_2(p_i, t_j) = \sum_{l=1}^9 s_{j,p_i}^{(l)} / \sum_{l=1}^9 O_{j,p_i}^{(l)}$ and $f_3(p_i, t_j) = \max_{l=1 \dots 9} s_{j,p_i}^{(l)}$, where $s_{j,p_i}^{(l)}$ is the anomaly score of signal l of product p_i at t_j . Finally, $f_4(p_i, t_j) = \sum_{l=1}^9 w_l \cdot s_{j,p_i}^{(l)}$ is the weighted sum of the anomaly scores, where $w_l = 1 / \sum_{t=1}^j O_{t,p_i}^{(l)}$. Simply put, the weight w_l is inverse proportional to the number of anomalies in signal l . The intuition is that the signals that contain a large number of anomalies may be noisy and provide us with many potential false positives.

We then use the empirical CDF to normalize the feature values, as shown in Equation 14.

$$F_g(p_i, t_j) = P(f_g \leq f_g(p_i, t_j)), \quad g = 1, \dots, 4 \quad (14)$$

where $F_g(p_i, t_j)$ is the ratio of f_g 's that are smaller than or equal to $f_g(p_i, t_j)$ across all products and all time points before t_j . The larger the $F_g(p_i, t_j)$, the larger the anomalousness. Finally, we compute p_i 's suspiciousness score $A(p_i, t_j)$ at time t_j as an average of the $F_g(p_i, t_j)$'s.

Experiments

Datasets

In this study we use review datasets from two different sites, namely SWM and FLIPKART, as we describe next.

SoftWare Marketplace (SWM) consists of reviews for all software products (apps) from the entertainment category (e.g., games, movies, sports, etc.) in an anonymous online marketplace which allows customers to purchase software applications. It contains 15,094 apps with over 1.1 million reviews by 966 thousand users, and spans 198 weeks between July 2008 and April 2012. The data for this market-

place was originally collected and used by (Akoglu, Chandy, and Faloutsos 2013).

FLIPKART contains reviews from flipkart.com, an e-commerce site which provides a platform for sellers to market products to customers. It contains 545 thousand products with roughly 3.3 million reviews by 1.1 million users, and spans 180 weeks between August 2011 to January 2015.

Results

Both the SWM and the FLIPKART datasets do not contain ground truth labels for anomalous (i.e., opinion-spammed) products, let alone the specific time periods in which each product was spammed. As such, we manually inspect the top-ranked products from both datasets, and provide evidence through various case studies.

SWM Case I: Game app We start with the further analysis and description of the product shown in Figure 1. This app, which started off with an average rating of 4, gradually declined below 3-star rating between weeks 75 to 165. A series of spam campaigns are then executed in weeks 168, 175, 182, and 189. Notice that these campaigns are organized every 7 weeks, which is a strong indication of manipulation.

When we use ‘Average Rating’ as the lead signal, we spot the first two campaigns, whereas when ‘No of + Reviews’ is used as the lead, all 4 weeks are detected (note the blue bars indicating the time points with anomaly score above the threshold). Nearly all the supporting signals also change simultaneously, suggesting low rating entropy and increased number of singleton reviewers.




Figure 2: Stacked bar charts showing daily review counts for the 4 detected campaigns in Figure 1 (week before, during, and after campaign separated by green vertical bars). Stacks represent counts for different ratings 1-5. Notice that spam campaigns involve mostly 5-star reviews (hence the bumpy increase in average rating week by week after each campaign).

Figure 2 shows the daily review counts for the week before, at, and after the spam campaigns, for each of the 4 detected campaigns, along with the distribution of ratings per day. Most of these reviews are from *singleton* reviewers that rated this app as 4 or 5 stars, as illustrated in Figure 3.

Next we investigate the review text to gain further insight into the campaigns. First we filter out all the negative (1-2 star) reviews of this app between weeks 75 to 165, during which the average rating was on a declining trend. Figure 4 (a) shows the word-cloud created from these reviews. It suggests that there is a ‘problem’ with the app as it ‘crashes’ often and needs a ‘fix’. Next we filter out the positive (4-5




Figure 3: Distribution of ratings by singleton reviewers only, during the 4 weeks of spam campaigns detected in Figure 1. This shows that the campaigns were carried by newly created accounts who provided a large number of 4- or 5-star ratings.

star) reviews from each of the spam campaign weeks, and show the resulting word-clouds in Figure 4 (b-e), respectively. We notice that the new reviewers think the app is “great” after the “new update” and they “love” it.

On further analysis we find that most of the reviews that singletons have written have duplicates or near duplicates. For example all the following text snippets appear multiple times across different reviews of this app: Great app for gamers, Great App For Gaming news, Great app, Easy to use, Must have app for gamers, One of my favorite apps, Use it every day, Very Informative.

SWM Case II: Movie app Our second case study is for a movie app from SWM. Figure 5 shows the timeline for the 9 indicative signals, where ‘Average Rating’ is designated as the lead. Week 149 is detected as anomalous at which the average rating increases from around 3.7 to above 4.

We explicitly show the anomaly scores (red curves) computed both by GLOBALAR (left) and LOCALAR (right) and the detected anomalous points (blue bars) for comparison. We notice that the anomalies detected by both methods are comparable, while LOCALAR computes scores only for those points around the anomalies indicated by the lead, being more efficient and as effective.

Figure 6 (top) shows the daily review counts, stacked by star-rating, before, during, and after the campaign week. This campaign is voluminous: a total of more than 4,000 reviews are written in week 149. The highest number of reviews per day is around 900 (May 21), which is roughly 1 review for every 2 minutes if the reviews were written by a single user. The piecharts in Figure 6 show the rating distribution of reviews from singletons (left) and non-singletons (right) during that week. Again, most spam reviews are 5-star rated and written from new accounts.

The following text is shared among singleton reviewers of this product at various weeks near the anomalous week: Best movie app ever!., Best movie app out there., Best movie app out there!, I use it all the time, I use this all the time!, I use

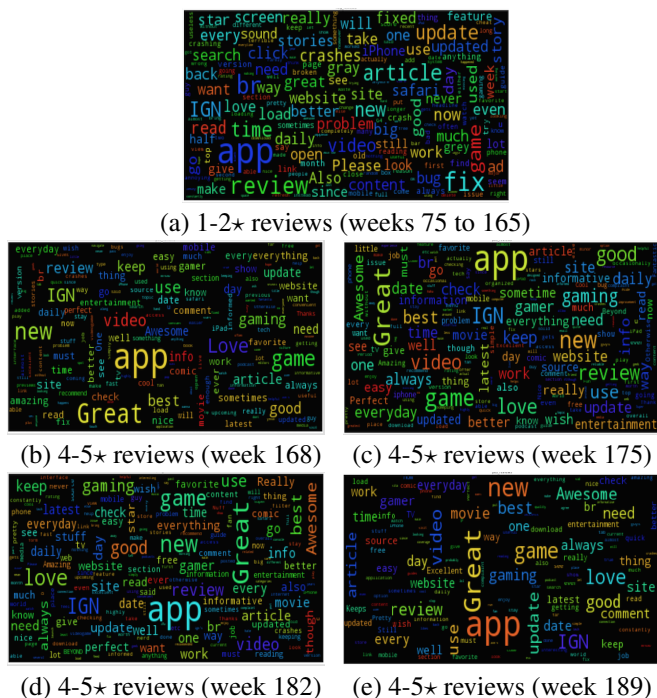


Figure 4: (a) Word cloud of negative reviews for time period during which average rating gradually declines for the product in Figure 1. (b-e) Word clouds of positive reviews during the respective 4 weeks of spam campaigns.

this app all the time., I use this app all the time!, Love it!, One of my favorite apps!!!, Love the app, use it all the time!, Way better than Fandango!.

SWM Case III: Whip-sound app Outliers in data need not always correspond to anomalies (in our case, spam campaigns) but also novelty or new trends. We present such an example from the SWM dataset.

The product “Pocket Whip” is a famous app that produces a whip sound. A TV show called the “Big Bang Theory” introduced this whip app in season 5 episode 19. It was aired on March 8, 2012 and consequently, reviewers started downloading the app and writing reviews from March 9, 2012 throughout the week. We detect this activity on week 191 in Figure 7 and characterize through Figure 8. Our approach is powerful enough not only to identify anomalous/suspicious review spam activities but also new trends: e.g., a new version of an app, a new chef in a restaurant, etc.

While the majority of new users liked the app (hence increase in average rating in week 191), number of negative reviews also increased (although not as much, see supporting signal in row 3 of Figure 7). This is also evident from Figure 9 that shows the daily rating distributions.

FLIPKART Case I In Figure 10 (left) we show one of the most suspicious products in FLIPKART, where week 35 is detected as anomalous. During this period, around 80 reviews are written. However the characteristics of this campaign is different in two aspects. First, most reviews are rated 3- or 4-stars, but only a few 5-stars (See Figure 11 (top)), while being able to increase average rating. These mixed ratings appear to be for better camouflage. Second, most reviewers are non-singletons (unlike in SWM) (notice

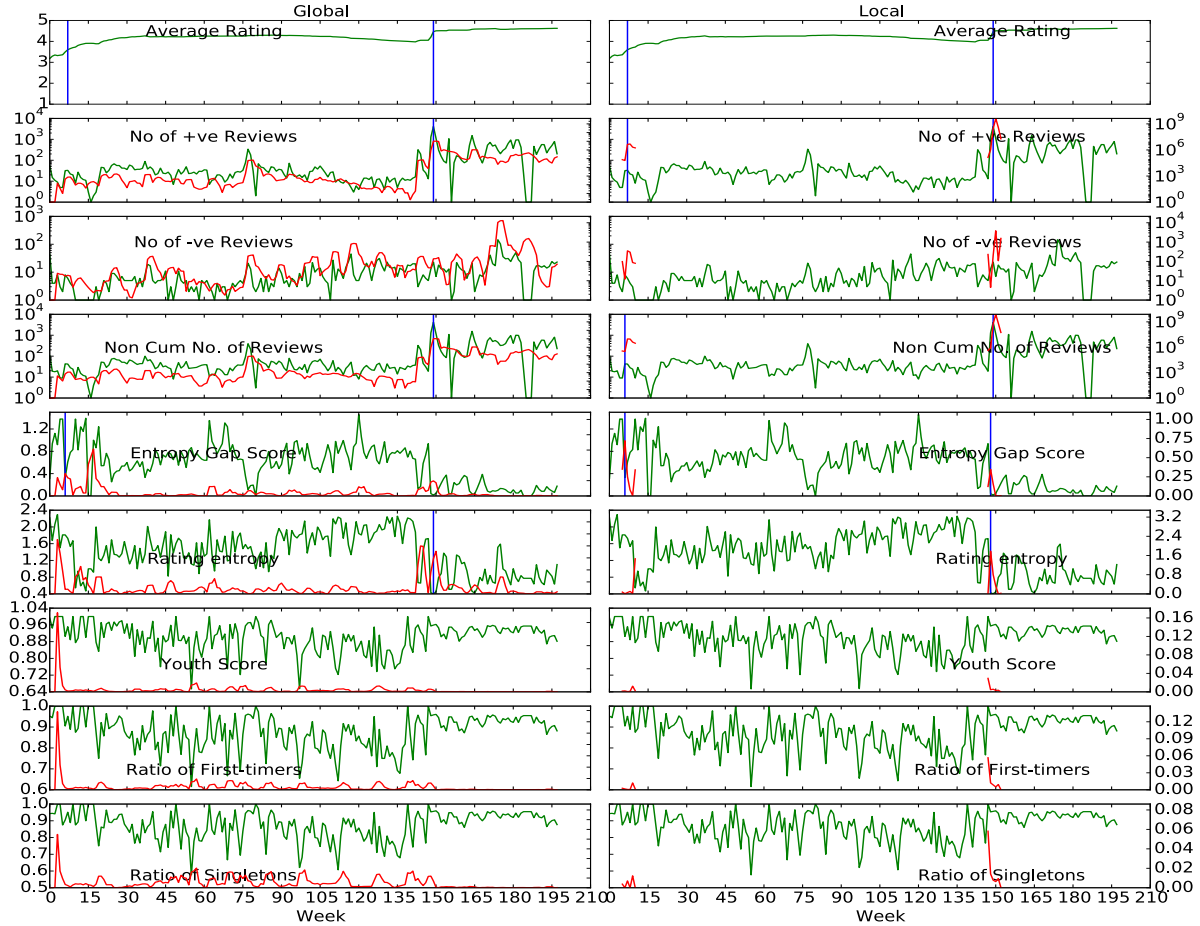


Figure 5: Time series for 9 indicative signals for a software product (app) in SWM (green curves with left y-axis: signal values, red curves with right y-axis: anomaly score, blue bars: anomalous time points). Lead signal: Average Rating (top row). Left: GLOBALAR results, Right: LOCALAR results. Note that while GLOBALAR computes anomaly scores for all time points for all signals, LOCALAR computes scores only for those time points indicated by the lead (notice the red curves), being efficient and equally effective.

no change in Ratio of First-timers and Singletons) although they are young accounts (notice spike in Youth Score). This suggests that other (related) products might also have been spammed by the same reviewers.

Figure 12 confirms this conjecture, where we find that a list of other products are rated similarly by these users *during the same time period*. Moreover, these periods are also detected as anomalous by our method. We show one other example product in Figure 10 (right) and its corresponding daily review count in Figure 11 (bottom)—notice the alignment in time and rating distribution (!) Further, we find that all these products are hair-related, including straighteners, dryers, and shavers, potentially belonging to the same seller.

FLIPKART Case II Our second case study is for an anomalous book from FLIPKART. As shown in Figure 13, its average rating increased to 4.4 on week 95. We find that this book received 125 5-star reviews in mainly two days during that week. Surprisingly those were from non-singletons (See Figure 14), who also reviewed another product—also a book (!) Further investigation revealed that those were 2 out of 3 books of an author. The one in Figure 13 had average rating 3.2 on Goodreads.com. Similarly, their other book

was rated 3.3 on Goodreads while its (potentially spammed) Flipkart rating is 4.5. We found that almost all 125 reviewers are *in common*, and write 5-star reviews 7 PM–11.45PM on June 8 and 11AM–7PM on June 9, 2013. What is more, their reviews follow nearly the *same order* for both books.

Conclusion

Opinion spam has become a prevalent problem, for which a vast body of methods operate in an offline fashion on a collection of static data. In this work, we brought emphasis to the aspect of time, and approached this problem with a novel temporal formulation. We proposed a new methodology that (1) monitors a comprehensive list of indicative signals over time, (2) spots anomalous events in real-time, and (3) provides descriptive pointers for manual inspection and characterization. As such, our approach exhibits desirable properties, as it is online, efficient, descriptive, and general.

Importantly, while we applied our methodology on opinion spam, it is general enough to be employed for other applications in which multiple signals are monitored over time, such as enterprise security, cyber-physical sensor systems, environmental monitoring, and surveillance systems.




Figure 6: (top) Daily review counts before, during, and after the spam campaign detected on week 149 in Figure 5. $>4,000$ reviews during campaign, highest review count per day: ≈ 900 . (bottom) Rating distribution of (majority of) reviews from singleton (left) and non-singleton reviewers (right) in week 149.

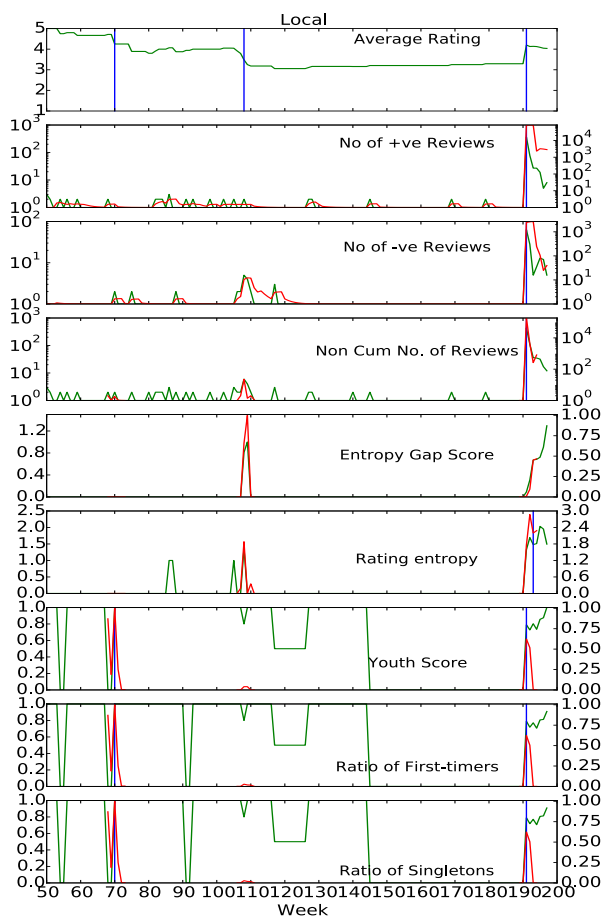


Figure 7: Time series for 9 indicative signals for a whip-sound app in SWM. On week 191, this app appeared in “Big Bang Theory” (a popular TV show) Season 5 Episode 19.



Figure 8: *Word cloud of reviews during week 191 in Figure 7.*




Figure 9: (top) Daily review counts before, during, and after the event detected on week 191 in Figure 7.

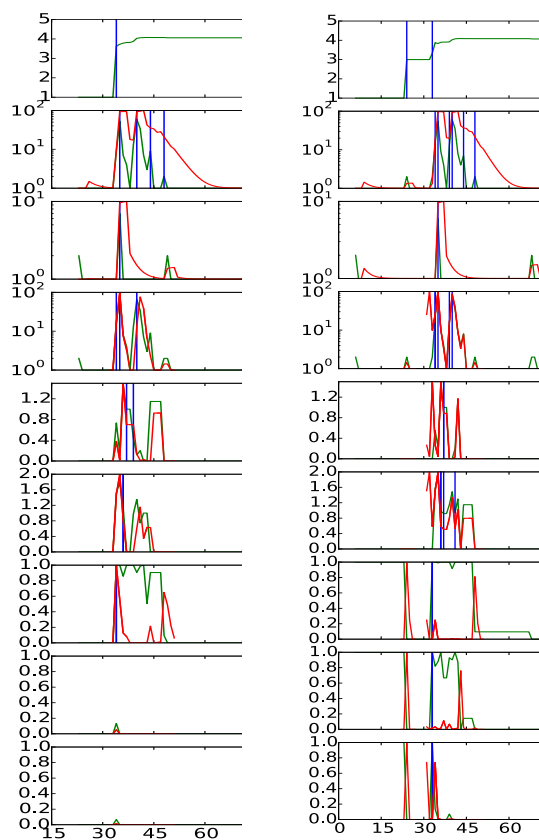


Figure 10: Partial time series for 9 indicative signals for two different products from FLIPKART that were spammed by same reviewers during same time periods (week 35 and 40).

References

- [Akoglu, Chandy, and Faloutsos 2013] Akoglu, L.; Chandy, R.; and Faloutsos, C. 2013. Opinion fraud detection in online reviews by network effects. In *ICWSM*.
- [Fei et al. 2013] Fei, G.; Mukherjee, A.; Liu, B.; Hsu, M.; Castel-




Figure 11: Daily review counts before, during, and after the spam campaign week 35 for two FLIPKART products in Figure 10.




Figure 12: Reviewers (left) concurrently spammed multiple hair products (right) from FLIPKART, including those in Figure 10. Edge labels: (rating, campaign week)



Figure 13: Time series for 3 lead signals for a FLIPKART book.

lanos, M.; and Ghosh, R. 2013. Exploiting burstiness in reviews for review spammer detection. In *ICWSM*.

[Feng, Banerjee, and Choi 2012] Feng, S.; Banerjee, R.; and Choi, Y. 2012. Syntactic stylometry for deception detection. In *ACL*.

[Feng et al. 2012] Feng, S.; Xing, L.; Gogar, A.; and Choi, Y. 2012. Distributional footprints of deceptive product reviews. *ICWSM* 12:98–105.

[Jindal and Liu 2008] Jindal, N., and Liu, B. 2008. Opinion spam and analysis. In *WSDM*, 219–230.

[Jindal, Liu, and Lim 2010] Jindal, N.; Liu, B.; and Lim, E.-P. 2010. Finding unusual review patterns using unexpected rules. In *CIKM*, 1549–1552.




Figure 14: (left) Daily review counts, and (right) Rating distribution by non-singletons in week 35 for product in Figure 13.

[Li et al. 2014] Li, H.; Chen, Z.; Liu, B.; Wei, X.; and Shao, J. 2014. Spotting fake reviews via collective positive-unlabeled learning. In *ICDM*, 899–904.

[Li et al. 2015] Li, H.; Chen, Z.; Mukherjee, A.; Liu, B.; and Shao, J. 2015. Analyzing and detecting opinion spam on a large-scale dataset via temporal and spatial patterns. In *ICWSM*, 634–637.

[Lim et al. 2010] Lim, E.-P.; Nguyen, V.-A.; Jindal, N.; Liu, B.; and Lauw, H. W. 2010. Detecting product review spammers using rating behaviors. In *CIKM*, 939–948.

[Luca 2011] Luca, M. 2011. Reviews, reputation, and revenue: The case of yelp.com. In *Working Paper 12-016, Harvard Bus. Sch.*

[Mukherjee et al. 2013a] Mukherjee, A.; Kumar, A.; Liu, B.; Wang, J.; Hsu, M.; Castellanos, M.; and Ghosh, R. 2013a. Spotting opinion spammers using behavioral footprints. In *KDD*.

[Mukherjee et al. 2013b] Mukherjee, A.; Venkataraman, V.; Liu, B.; and Glance, N. S. 2013b. What yelp fake review filter might be doing? In *ICWSM*.

[Mukherjee, Liu, and Glance 2012] Mukherjee, A.; Liu, B.; and Glance, N. S. 2012. Spotting fake reviewer groups in consumer reviews. In *WWW*, 191–200.

[Newman 2005] Newman, M. 2005. Power laws, Pareto distributions and Zipf’s law. *Contemporary Physics* 46(5):323–351.

[Ott et al. 2011] Ott, M.; Choi, Y.; Cardie, C.; and Hancock, J. T. 2011. Finding deceptive opinion spam by any stretch of the imagination. In *ACL*, 309–319.

[Page 1954] Page, E. S. 1954. Continuous Inspection Schemes. *Biometrika* 41(1/2):100–115.

[Rahman et al. 2015] Rahman, M.; Carbanar, B.; Ballesteros, J.; and Chau, D. H. 2015. To catch a fake: Curbing deceptive yelp ratings and venues. *Stat. Analysis and Data Min.* 8(3):147–161.

[Rayana and Akoglu 2015] Rayana, S., and Akoglu, L. 2015. Collective opinion spam detection: Bridging review networks and metadata. In *KDD*, 985–994.

[Smets and Vreeken 2011] Smets, K., and Vreeken, J. 2011. The odd one out: Identifying and characterising anomalies. In *SDM*, volume 108, 109–148. SIAM.

[Streitfeld 2011] Streitfeld, D. 2011. In a Race to Out-Rave, 5-Star Web Reviews Go for 5 Dollars. <http://nyti.ms/1l9wqXW>.

[Wang et al. 2011] Wang, G.; Xie, S.; Liu, B.; and Yu, P. S. 2011. Review graph based online store review spammer detection. In *ICDM*, 1242–1247.

[Xie et al. 2012] Xie, S.; Wang, G.; Lin, S.; and Yu, P. S. 2012. Review spam detection via temporal pattern discovery. In *KDD*, 823–831.

[Yamanishi and Takeuchi 2002] Yamanishi, K., and Takeuchi, J.-i. 2002. A unifying framework for detecting outliers and change points from non-stationary time series data. In *KDD*, 676–681.

[Ye and Akoglu 2015] Ye, J., and Akoglu, L. 2015. Discovering opinion spammer groups by network footprints. In *ECML/PKDD*, 267–282.